

<b>UTILITY PATENT APPLICATION TRANSMITTAL</b> <small>(Only for new nonprovisional applications under 37 CFR 1.53(b))</small>	Attorney Docket No.	SJ00-00-044	Total Pages	37
	First Named Inventor or Application Identifier			
	Kevin Frank Smith			
	Express Mail Label No.	EJ137621500US		

<b>APPLICATION ELEMENTS</b> <i>See MPEP chapter 600 concerning utility patent application contents.</i>	<b>ADDRESS TO:</b> Assistant Commissioner for Patents Box Patent Application Washington, DC 20231	
1. <input checked="" type="checkbox"/> Fee Transmittal Form <i>(Submit an original, and a duplicate for fee processing)</i>	6. <input type="checkbox"/> Microfiche Computer Program <i>(Appendix)</i>	
2. <input checked="" type="checkbox"/> Specification <i>(preferred arrangement set forth below)</i> [Total Pages] <b>30</b> <div>- Descriptive title of the invention - Cross References to Related Application - Statement Regarding Fed sponsored R &amp; D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings <i>(if filed)</i> - Detailed Description - Claim(s) - Abstract of the Disclosure</div>	7. Nucleotide and/or Amino Acid Sequence Submission <i>(if applicable, all necessary)</i> <div>a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy <i>(identical to computer copy)</i> c. <input type="checkbox"/> Statement verifying identify of above copies</div>	
3. <input checked="" type="checkbox"/> Drawing(s) <i>(35 USC 113)</i> [Total Sheets] <b>5</b>	<b>ACCOMPANYING APPLICATION PARTS</b>	
4. Oath or Declaration [Total Pages] <b>2</b> <div>a. <input checked="" type="checkbox"/> Newly executed <i>(original or copy)</i> b. <input type="checkbox"/> Copy from a prior application <i>(37 CFR 1.63(d))</i> <i>(for continuation /divisional with Box 17 completed)</i> <b>[Note Box 5 below]</b> i. <input type="checkbox"/> <b>DELETION OF INVENTOR(S)</b> Signed statement attached deleting inventor(s) named in prior application, see 37 CFR 1.63(d)(2) and 1.33(b).</div>	8. <input checked="" type="checkbox"/> Assignment Papers <i>(cover sheet &amp; document(s))</i>	
5. <input type="checkbox"/> Incorporation by Reference <i>(useable if Box 4b is checked)</i> The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.	9. <input type="checkbox"/> 37 CFR 3.73(b) Statement <input type="checkbox"/> Power of Attorney <i>(when there is an assignee)</i>	
	10. <input type="checkbox"/> English Translation Document <i>(if applicable)</i>	
	11. <input checked="" type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input checked="" type="checkbox"/> Copies of IDS Citations	
	12. <input type="checkbox"/> Preliminary Amendment	
	13. <input checked="" type="checkbox"/> Return Receipt Postcard <i>(MPEP 503)</i> <i>(Should be specifically itemized)</i>	
	14. <input type="checkbox"/> Small Entity <input type="checkbox"/> Statement filed in prior application, Statement(s) Status still proper and desired	
	15. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>if foreign priority is claimed)</i>	
	16. <input checked="" type="checkbox"/> Other: Express Mail Certificate	

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:

<input type="checkbox"/> Continuation	<input type="checkbox"/> Divisional	<input type="checkbox"/> Continuation-in-part (CIP)	of prior application No.:
---------------------------------------	-------------------------------------	---	---------------------------

18. **CORRESPONDENCE ADDRESS**

☐ Customer Number or Bar Code Label or ☐ Correspondence address below

NAME	Brian C. Kunzler				
ADDRESS	10 West 100 South				
CITY	Salt Lake City	STATE	Utah	ZIP	84101
COUNTRY	United States	TELEPHONE	(801) 994-4646	FAX	(801) 322-1054

Please type a plus sign (+) inside this box → +

PTO/SB/05 (6)  
Approved for use through 9/30/00 OMB 0651-0046  
Patent and Trademark Office, U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

<h2 style="margin: 0;">FEE TRANSMITTAL</h2> <p style="font-size: small; margin-top: 10px;">Note: Effective October 1, 1997. Patent fees are subject to annual revision.</p>		<b>Complete If Known</b>	
		Application Number	Not yet assigned
		Filing Date	October 12, 2000
		First Named Inventor	Kevin Frank Smith
		Group Art Unit	
		Examiner Name	
TOTAL AMOUNT OF PAYMENT	\$ 894.00	Attorney Docket Number	SJO0-00-044

METHOD OF PAYMENT (check one)	FEE CALCULATION (continued)																																																																																																																																																																						
<p>1. <input checked="" type="checkbox"/> The Commissioner is hereby authorized to charge indicated fees and credit any over payments to:</p> <p>Deposit Account Number: <u>09-0446</u></p> <p>Deposit Account Name: <u>IBM CORPORATION</u></p> <p><input checked="" type="checkbox"/> Charge Any Additional Fee Required Under 37 CFR 1.16 and 1.17      <input type="checkbox"/> Charge the Issue Fee In 37 CFR at the Mailing of the Notice of Allowance</p> <p>2. <input type="checkbox"/> Payment Enclosed:</p> <p style="margin-left: 20px;"><input type="checkbox"/> Check    <input type="checkbox"/> Money Order    <input type="checkbox"/> Other</p>	<p><b>3. ADDITIONAL FEES</b></p> <table style="width: 100%; font-size: small;"> <thead> <tr> <th colspan="2">Large Entity</th> <th colspan="2">Small Entity</th> <th rowspan="2">Fee Description</th> <th rowspan="2">Fee Paid</th> </tr> <tr> <th>Fee Code</th> <th>Fee (\$)</th> <th>Fee Code</th> <th>Fee (\$)</th> </tr> </thead> <tbody> <tr><td>105</td><td>130</td><td>205</td><td>65</td><td>Surcharge - late filing fee or oath</td><td></td></tr> <tr><td>127</td><td>50</td><td>227</td><td>25</td><td>Surcharge - late provisional filing fee or cover sheet</td><td></td></tr> <tr><td>139</td><td>130</td><td>139</td><td>130</td><td>Non-English specification</td><td></td></tr> <tr><td>147</td><td>2520</td><td>147</td><td>2520</td><td>For filing a request for reexamination</td><td></td></tr> <tr><td>112</td><td>920*</td><td>112</td><td>920*</td><td>Requesting publication of SIR prior to Examiner action</td><td></td></tr> <tr><td>113</td><td>1840*</td><td>113</td><td>1840*</td><td>Requesting publication of SIR after Examiner action</td><td></td></tr> <tr><td>115</td><td>110</td><td>215</td><td>55</td><td>Extension for reply within first month</td><td></td></tr> <tr><td>116</td><td>390</td><td>216</td><td>195</td><td>Extension for reply within second month</td><td></td></tr> <tr><td>117</td><td>890</td><td>217</td><td>445</td><td>Extension for reply within third month</td><td></td></tr> <tr><td>118</td><td>1390</td><td>218</td><td>695</td><td>Extension for reply within fourth month</td><td></td></tr> <tr><td>128</td><td>1890</td><td>228</td><td>945</td><td>Extension for reply within fifth month</td><td></td></tr> <tr><td>119</td><td>310</td><td>219</td><td>155</td><td>Notice of Appeal</td><td></td></tr> <tr><td>120</td><td>310</td><td>220</td><td>155</td><td>Filing a brief in support of an appeal</td><td></td></tr> <tr><td>121</td><td>270</td><td>221</td><td>135</td><td>Request for oral hearing</td><td></td></tr> <tr><td>138</td><td>1510</td><td>138</td><td>1510</td><td>Petition to institute a public use proceeding</td><td></td></tr> <tr><td>140</td><td>110</td><td>240</td><td>55</td><td>Petition to revive - intentional</td><td></td></tr> <tr><td>141</td><td>1240</td><td>241</td><td>620</td><td>Petition to revive - unintentional</td><td></td></tr> <tr><td>142</td><td>1210</td><td>242</td><td>605</td><td>Utility issue fee</td><td></td></tr> <tr><td>143</td><td>430</td><td>243</td><td>215</td><td>Design issue fee</td><td></td></tr> <tr><td>144</td><td>580</td><td>244</td><td>290</td><td>Plant issue fee</td><td></td></tr> <tr><td>122</td><td>130</td><td>122</td><td>130</td><td>Petitions to the Commissioner</td><td></td></tr> <tr><td>123</td><td>50</td><td>123</td><td>50</td><td>Petitions related to provisional applications</td><td></td></tr> <tr><td>126</td><td>240</td><td>126</td><td>240</td><td>Submission of Information Disclosure Stmt</td><td></td></tr> <tr><td>581</td><td>40</td><td>581</td><td>40</td><td>Recording each patent assignment per property (times number of properties)</td><td>40</td></tr> <tr><td>146</td><td>760</td><td>246</td><td>380</td><td>Filing a submission after final rejection (37 CFR 1.129(a))</td><td></td></tr> <tr><td>149</td><td>760</td><td>249</td><td>380</td><td>For each additional invention to be examined (37 CFR 1.129(b))</td><td></td></tr> </tbody> </table> <p>Other fee (specify) _____</p> <p>Other fee (specify) _____</p>	Large Entity		Small Entity		Fee Description	Fee Paid	Fee Code	Fee (\$)	Fee Code	Fee (\$)	105	130	205	65	Surcharge - late filing fee or oath		127	50	227	25	Surcharge - late provisional filing fee or cover sheet		139	130	139	130	Non-English specification		147	2520	147	2520	For filing a request for reexamination		112	920*	112	920*	Requesting publication of SIR prior to Examiner action		113	1840*	113	1840*	Requesting publication of SIR after Examiner action		115	110	215	55	Extension for reply within first month		116	390	216	195	Extension for reply within second month		117	890	217	445	Extension for reply within third month		118	1390	218	695	Extension for reply within fourth month		128	1890	228	945	Extension for reply within fifth month		119	310	219	155	Notice of Appeal		120	310	220	155	Filing a brief in support of an appeal		121	270	221	135	Request for oral hearing		138	1510	138	1510	Petition to institute a public use proceeding		140	110	240	55	Petition to revive - intentional		141	1240	241	620	Petition to revive - unintentional		142	1210	242	605	Utility issue fee		143	430	243	215	Design issue fee		144	580	244	290	Plant issue fee		122	130	122	130	Petitions to the Commissioner		123	50	123	50	Petitions related to provisional applications		126	240	126	240	Submission of Information Disclosure Stmt		581	40	581	40	Recording each patent assignment per property (times number of properties)	40	146	760	246	380	Filing a submission after final rejection (37 CFR 1.129(a))		149	760	249	380	For each additional invention to be examined (37 CFR 1.129(b))	
Large Entity		Small Entity		Fee Description	Fee Paid																																																																																																																																																																		
Fee Code	Fee (\$)	Fee Code	Fee (\$)																																																																																																																																																																				
105	130	205	65	Surcharge - late filing fee or oath																																																																																																																																																																			
127	50	227	25	Surcharge - late provisional filing fee or cover sheet																																																																																																																																																																			
139	130	139	130	Non-English specification																																																																																																																																																																			
147	2520	147	2520	For filing a request for reexamination																																																																																																																																																																			
112	920*	112	920*	Requesting publication of SIR prior to Examiner action																																																																																																																																																																			
113	1840*	113	1840*	Requesting publication of SIR after Examiner action																																																																																																																																																																			
115	110	215	55	Extension for reply within first month																																																																																																																																																																			
116	390	216	195	Extension for reply within second month																																																																																																																																																																			
117	890	217	445	Extension for reply within third month																																																																																																																																																																			
118	1390	218	695	Extension for reply within fourth month																																																																																																																																																																			
128	1890	228	945	Extension for reply within fifth month																																																																																																																																																																			
119	310	219	155	Notice of Appeal																																																																																																																																																																			
120	310	220	155	Filing a brief in support of an appeal																																																																																																																																																																			
121	270	221	135	Request for oral hearing																																																																																																																																																																			
138	1510	138	1510	Petition to institute a public use proceeding																																																																																																																																																																			
140	110	240	55	Petition to revive - intentional																																																																																																																																																																			
141	1240	241	620	Petition to revive - unintentional																																																																																																																																																																			
142	1210	242	605	Utility issue fee																																																																																																																																																																			
143	430	243	215	Design issue fee																																																																																																																																																																			
144	580	244	290	Plant issue fee																																																																																																																																																																			
122	130	122	130	Petitions to the Commissioner																																																																																																																																																																			
123	50	123	50	Petitions related to provisional applications																																																																																																																																																																			
126	240	126	240	Submission of Information Disclosure Stmt																																																																																																																																																																			
581	40	581	40	Recording each patent assignment per property (times number of properties)	40																																																																																																																																																																		
146	760	246	380	Filing a submission after final rejection (37 CFR 1.129(a))																																																																																																																																																																			
149	760	249	380	For each additional invention to be examined (37 CFR 1.129(b))																																																																																																																																																																			
<p><b>FEE CALCULATION</b></p> <p><b>1. FILING FEE</b></p> <table style="width: 100%; font-size: small;"> <thead> <tr> <th colspan="2">Large Entity</th> <th colspan="2">Small Entity</th> <th rowspan="2">Fee Description</th> <th rowspan="2">Fee Paid</th> </tr> <tr> <th>Fee Code</th> <th>Fee (\$)</th> <th>Fee Code</th> <th>Fee (\$)</th> </tr> </thead> <tbody> <tr><td>101</td><td>710</td><td>201</td><td>355</td><td>Utility filing fee</td><td>710</td></tr> <tr><td>106</td><td>320</td><td>206</td><td>160</td><td>Design filing fee</td><td></td></tr> <tr><td>107</td><td>490</td><td>207</td><td>245</td><td>Plant filing fee</td><td></td></tr> <tr><td>108</td><td>710</td><td>208</td><td>355</td><td>Reissue filing fee</td><td></td></tr> <tr><td>114</td><td>150</td><td>214</td><td>75</td><td>Provisional filing fee</td><td></td></tr> <tr><td colspan="5"><b>SUBTOTAL (1)</b></td><td><b>\$ 710</b></td></tr> </tbody> </table> <p><b>2. CLAIMS</b></p> <table style="width: 100%; font-size: small;"> <thead> <tr> <th></th> <th>Extra</th> <th>Fee from below</th> <th>Fee Paid</th> </tr> </thead> <tbody> <tr> <td>Total Claims    28    -20 =</td> <td>8</td> <td>18</td> <td>144</td> </tr> <tr> <td>Ind. Claims    3    -3 =</td> <td>0</td> <td>80</td> <td>0</td> </tr> <tr> <td>Multiple Dep. Claims</td> <td>0</td> <td>270</td> <td>0</td> </tr> </tbody> </table> <table style="width: 100%; font-size: small;"> <thead> <tr> <th colspan="2">Large Entity</th> <th colspan="2">Small Entity</th> <th rowspan="2">Fee Description</th> <th rowspan="2">Fee Paid</th> </tr> <tr> <th>Fee Code</th> <th>Fee (\$)</th> <th>Fee Code</th> <th>Fee (\$)</th> </tr> </thead> <tbody> <tr><td>103</td><td>18</td><td>203</td><td>9</td><td>Claims in excess of 20</td><td></td></tr> <tr><td>102</td><td>80</td><td>202</td><td>40</td><td>Independent claims in excess of 3</td><td></td></tr> <tr><td>104</td><td>270</td><td>204</td><td>135</td><td>Multiple dependent claim</td><td></td></tr> <tr><td>109</td><td>80</td><td>209</td><td>40</td><td>Reissue independent claims over original patent</td><td></td></tr> <tr><td>110</td><td>18</td><td>210</td><td>9</td><td>Reissue claims in excess of 20 and over original patent</td><td></td></tr> <tr><td colspan="5"><b>SUBTOTAL (2)</b></td><td><b>144</b></td></tr> </tbody> </table>	Large Entity		Small Entity		Fee Description	Fee Paid	Fee Code	Fee (\$)	Fee Code	Fee (\$)	101	710	201	355	Utility filing fee	710	106	320	206	160	Design filing fee		107	490	207	245	Plant filing fee		108	710	208	355	Reissue filing fee		114	150	214	75	Provisional filing fee		<b>SUBTOTAL (1)</b>					<b>\$ 710</b>		Extra	Fee from below	Fee Paid	Total Claims    28    -20 =	8	18	144	Ind. Claims    3    -3 =	0	80	0	Multiple Dep. Claims	0	270	0	Large Entity		Small Entity		Fee Description	Fee Paid	Fee Code	Fee (\$)	Fee Code	Fee (\$)	103	18	203	9	Claims in excess of 20		102	80	202	40	Independent claims in excess of 3		104	270	204	135	Multiple dependent claim		109	80	209	40	Reissue independent claims over original patent		110	18	210	9	Reissue claims in excess of 20 and over original patent		<b>SUBTOTAL (2)</b>					<b>144</b>	<p><b>SUBTOTAL (3)</b>      <b>40</b></p> <p style="font-size: x-small;">*Reduced by Basic Filing Fee</p>																																																										
Large Entity		Small Entity		Fee Description			Fee Paid																																																																																																																																																																
Fee Code	Fee (\$)	Fee Code	Fee (\$)																																																																																																																																																																				
101	710	201	355	Utility filing fee	710																																																																																																																																																																		
106	320	206	160	Design filing fee																																																																																																																																																																			
107	490	207	245	Plant filing fee																																																																																																																																																																			
108	710	208	355	Reissue filing fee																																																																																																																																																																			
114	150	214	75	Provisional filing fee																																																																																																																																																																			
<b>SUBTOTAL (1)</b>					<b>\$ 710</b>																																																																																																																																																																		
	Extra	Fee from below	Fee Paid																																																																																																																																																																				
Total Claims    28    -20 =	8	18	144																																																																																																																																																																				
Ind. Claims    3    -3 =	0	80	0																																																																																																																																																																				
Multiple Dep. Claims	0	270	0																																																																																																																																																																				
Large Entity		Small Entity		Fee Description	Fee Paid																																																																																																																																																																		
Fee Code	Fee (\$)	Fee Code	Fee (\$)																																																																																																																																																																				
103	18	203	9	Claims in excess of 20																																																																																																																																																																			
102	80	202	40	Independent claims in excess of 3																																																																																																																																																																			
104	270	204	135	Multiple dependent claim																																																																																																																																																																			
109	80	209	40	Reissue independent claims over original patent																																																																																																																																																																			
110	18	210	9	Reissue claims in excess of 20 and over original patent																																																																																																																																																																			
<b>SUBTOTAL (2)</b>					<b>144</b>																																																																																																																																																																		

SUBMITTED BY				Complete (if applicable)	
Typed or Printed Name	Brian C. Kunzler			Reg. Number	38,527
Signature		Date	Oct 12, 2000	Deposit Account User ID	

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Washington, DC 20231.

PTO/SB/05 (6)  
 10/12/00  
 109/009188

1                                    **SCHEDULING SEQUENTIAL DATA**  
2                                    **PREFETCHES IN A PREEXISTING LRU CACHE**  
3

4  
5                                    **BACKGROUND OF THE INVENTION**

6                    1.        **The Field of the Invention**  
7

8                    The present invention relates to cache management in data storage systems. More  
9 specifically, the present invention relates to prefetch scheduling in a preexisting LRU cache  
10 of a data storage system.  
11

12  
13                   2.        **The Relevant Art**

14                   Cache memory is used in data storage systems to buffer frequently accessed data in  
15 order to allow the data to be accessed at a relatively high rate. The cache memory is a  
16 relatively small high speed memory operating on a host processor or between the host  
17 processor and relatively slower memory devices. Typical data storage systems using caching  
18 may include a cache directory or index of the data elements in a main memory of the hosts  
19 operating on the data storage system. The cache directory is referenced to provide an  
20 indication of whether or not each data element of the main memory resides in the cache  
21 memory at any give time, and if so, to indicate the present location of the data element in the  
22 cache memory. When a host processor requests an Input/Output (I/O) operation, the cache  
23 directory is first consulted to determine whether the requested data element is present in the  
24 cache memory and if so, to determine its location. When the data element is present in the  
25  
26

1 cache memory, the data element can be quickly accessed, rather than having to be requested  
2 from a slower storage device.

3  
4 Generally, in such systems, every time a data element is requested, a determination  
5 is made whether the accessed data is likely to be accessed again in the near future. If so, the  
6 accessed data element is copied or “staged” into the cache memory. In some data storage  
7 systems, requested data elements are always staged into the cache memory if they are absent  
8 from the cache memory. Some data storage systems are also responsive to explicit “prefetch”  
9 commands from the host computer to cause specified data to be staged into the cache, even  
10 though the specified data is not immediately accessed by the host computer.  
11  
12

13 Because the cache memory has a capacity that is smaller than the main memory, it  
14 is frequently necessary for data elements in the cache memory to be replaced or removed from  
15 the cache memory in order to provide space in the cache memory for more recently requested  
16 data elements. In general, for the cache memory to be useful, the data elements removed or  
17 replaced from the cache memory must be calculated to be less likely to be accessed in the near  
18 future than the new data elements being staged into the cache memory at the time the removal  
19 or replacement occurs.  
20  
21

22 Data storage systems that use disk drives for the main memory typically use random  
23 access memory (RAM) for the cache memory. In such a data storage system, the data  
24 elements in the cache memory are often logical tracks of data on the disks, although in many  
25 systems, the data records are blocks or records of data. The cache directory includes a  
26 directory entry for at least each data element stored in the cache. Each directory entry for  
each data element stored in the cache memory generally includes a pointer to the location of

1 the data element in the cache memory. The cache directory can be a table including an entry  
2 for each data element stored in the disk storage. Alternatively, the directory may include a  
3 hash table for accessing lists of the directory entries so that the cache directory need not  
4 include any cache directory entries for data elements that are absent from the cache memory.  
5 In either case, any one of a plurality of data elements in the cache memory may be replaced  
6 or removed from the cache according to the particular cache management scheme being used  
7 to make room for another data element.  
8

10 The performance of such a data storage system is highly dependent on the cache  
11 management scheme used for selecting the data element to be removed or replaced. The  
12 cache management scheme is implemented by a cache management system, or “cache  
13 manager,” in the data storage system.  
14

15 In one common cache management scheme, a cache manager is programmed to  
16 remove or replace the “least-recently-used” (LRU) data element in the cache memory. The  
17 least-recently-used data element is usually the data element accessed least recently by the host  
18 computer. The cache manager maintains an ordered list, or queue, of the data elements in  
19 the cache memory so that the cache manager can readily identify the least-recently-used data  
20 element. The queue is typically maintained in a doubly-linked list. When a data element is  
21 accessed, the data element is moved to the head of the queue, unless the data element is  
22 already at the head of the queue. This process is known as making the data element “young”  
23 in the cache and ensures that, when the queue is not empty, the least-recently-used data  
24 element in the cache memory will be located at the end of the queue and the most-recently-  
25 used element in the cache memory will be located at the head of the queue.  
26

1 Data that is fetched into the cache memory may be described in two broad fashions.  
2 The first is random access data which denotes data that is needed for specific operations but  
3 which is not connected with other data in any manner. Many caching systems are configured  
4 for optimal performance when fetching random access data. The second type of data is  
5 known as sequential data, denoting that several elements of the data are used by a processor  
6 in a specific sequence, typically the sequence in which the data elements are stored on a  
7 storage device. Many systems that employ a dedicated or “native” LRU cache are only  
8 designed to store data accessed in random access operations and make no provision for  
9 accessing sequential data.  
10  
11

12 Attempts have been made to improve the performance of a native LRU cache when  
13 fetching sequential data. These solutions, however, require a modification in some fashion  
14 of the LRU cache itself to achieve satisfactory performance for sequential data prefetches.  
15 One example of these types of modifications is the creation of a “microcache” within the  
16 existing LRU cache to hold the sequential data.  
17  
18

19 Modifying existing LRU caches is not always a plausible solution. For instance, in  
20 legacy systems it may not be possible or desirable to modify the replacement algorithm of the  
21 cache. The cache logic or controller may be inaccessible or hardwired, or the system the  
22 cache resides on may have been provided for a specific purpose that modifying the cache  
23 would disrupt.  
24  
25

26 Accordingly, a need exists in the art for a method of scheduling prefetches of  
sequential data into a native LRU cache without directly modifying the algorithm or structure  
of the LRU cache.



1 the number of logically preceding data elements present in the cache according to the model  
2 of the cache. The assigned priority value is compared against a threshold value to determine  
3 if the requested data element is to be prefetched. If the value assigned to the data requested  
4 is greater than the threshold value , the prefetch scheduling system schedules one or more  
5 prefetches of logically successive data elements.  
6

7  
8 Scheduling a prefetch may comprise sending an Input/Output (I/O) request for the  
9 logically successive data element to the preexisting LRU cache, which then loads the  
10 successive data element into the cache. Because the successive data element was unsolicited  
11 by the host processor, the data element is ignored by the host processor.  
12

13 The threshold value is preferably dynamic and is adjusted periodically as needed,  
14 according to the history of prefetched data elements within the cache. If prefetched data  
15 elements have a history of being hit more than other data elements within the cache, the  
16 threshold value is decremented, and if prefetched data elements are falling out of the cache  
17 without being hit at a greater rate than other elements, the threshold value is incremented.  
18  
19  
20  
21  
22  
23  
24  
25  
26



## **BRIEF DESCRIPTION OF THE DRAWINGS**

In order that the manner in which the advantages and objects of the invention are obtained will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a schematic block diagram of a computer system suitable for implementing certain embodiments of the present invention.

Figure 2 is a schematic block diagram illustrating one embodiment of a data prefetch module of the present invention.

Figure 3 is a schematic flow chart diagram illustrating one embodiment of a data prefetch scheduling method of the present invention.

Figure 4 is a schematic flow chart diagram illustrating one embodiment of a data prefetch determination of the method of Figure 3.

Figure 5 is a schematic flow chart diagram illustrating one embodiment of a method for determining and updating a threshold value of the present invention.

## **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

Figure 1 is a schematic block diagram illustrating a computer system 10 in which executable and operational data, operating in accordance with the present invention, may be hosted on one or more computer stations 12 in a network 14. The network 14 preferably comprises a storage area network (SAN) but may also comprise a wide area network (WAN) or local area network (LAN) and may also comprise an interconnected system of networks, one particular example of which is the Internet and the World Wide Web supported on the Internet.

A typical computer station 12 may include a processor or CPU 16. The CPU 16 may be operably connected to one or more memory devices 18. The memory devices 18 are depicted as including a non-volatile storage device 20 such as a hard disk drive or CD-ROM drive, a read-only memory (ROM) 22, and a random access volatile memory (RAM) 24. Preferably, the computer station 12 operates under the control of an operating system (OS) 25, such as MVS®, OS/390®, AIX®, OS/2®, WINDOWS NT®, WINDOWS®, UNIX®, and the like.

The computer station 12 or system 10 in general may also include one or more input devices 26, such as a mouse or keyboard, for receiving inputs from a user or from another device. Similarly, one or more output devices 28, such as a monitor or printer, may be provided within or be accessible from the computer system 10. A network port such as a network interface card 30 may be provided for connecting to outside devices through the network 14. In the case where the network 14 is remote from the computer station, the

1 network interface card 30 may comprise a modem, and may connect to the network 14  
2 through a local access line such as a telephone line.

3  
4 Within any given station 12, a system bus 32 may operably interconnect the CPU 16,  
5 the memory devices 18, the input devices 26, the output devices 28, the network card 30, and  
6 one or more additional ports 34. The system bus 32 and a network backbone 36 may be  
7 regarded as data carriers. As such, the system bus 32 and the network backbone 36 may be  
8 embodied in numerous configurations. For instance, wire, fiber optic line, wireless  
9 electromagnetic communications by visible light, infrared, and radio frequencies may be  
10 implemented as appropriate.  
11

12  
13 In general, the network 14 may comprise a storage area network (SAN), local area  
14 network (LAN), a wide area network (WAN), several adjoining networks, an Intranet, or as  
15 in the manner depicted, a system of interconnected networks such as the Internet 40. The  
16 individual stations 12 communicate with each other over the backbone 36 and/or over the  
17 Internet 40 with varying degrees and types of communication capabilities and logic capability.  
18 The individual stations 12 may include a mainframe computer on which the modules of the  
19 present invention may be hosted.  
20  
21

22 Different communication protocols, e.g., fiber channel, ISO/OSI, IPX, TCP/IP, may  
23 be used on the network, but in the case of the Internet, a single, layered communications  
24 protocol (TCP/IP) generally enables communications between the differing networks 14 and  
25 stations 12. Thus, a communication link may exist, in general, between any of the stations  
26  
12.

1           The stations 12 connected on the network 14 may comprise data storage servers 46  
2 and/or data storage devices 45 to which may be connected an existing least-recently-used  
3 (LRU) cache 42. In the depicted embodiment, the cache 42 is a stand-alone module remote  
4 to both the station 12 and the server 46, but, of course, could be implemented within a station  
5 12 or a server 46, including a storage server. Other resources or peripherals 44, such as  
6 printers and scanners may also be connected to the network 36. Other networks may be in  
7 communication with the network 14 through a router 38 and/or over the Internet 40.  
8  
9

10           Figure 2 is a schematic block diagram illustrating one embodiment of a prefetch  
11 module 200, suitable for use in a data prefetch scheduling system. In accordance with the  
12 invention, the data prefetch scheduling system may also include a station 12, a cache 42, and  
13 a server 46 of Figure 1. In one embodiment, sequential data stored in a storage device 45 of  
14 the server 46 is prestaged into the cache 42 of Figure 1 using the prefetch module 200. The  
15 prefetch module 200 may reside anywhere in the computer system 10, and in a preferred  
16 embodiment resides as a daemon on the computer station 12 of Figure 1. More preferably,  
17 however, the prefetch module 200 is located on or between the computer station 12 (the  
18 "host") and the data storage server 46 of Figure 1, and most preferably, the prefetch module  
19 200 operates on a processor of the computer station 12. The prefetch module 200, as  
20 depicted, is configured with an interface module 202, a calculation module 204, a dynamic  
21 threshold optimization module 206, a prefetch request module 208, and a remote modeling  
22 module 210.  
23  
24  
25  
26

          The remote modeling module 210 preferably models the operation, and to the extent  
known, the contents of the cache 42 of Figure 1. The data prefetch module 200 preferably

1 uses the information from the modeling module 210 to schedule prefetches of sequential data  
2 from the main memory of the data storage server 46 into the LRU cache 42. The data  
3 prefetch module 200 maintains one or more models 220 of the LRU cache 42 of Figure 1 and  
4 in conjunction with those models 220, stores information about each data element in the cache  
5 as objects 230.  
6

7 Each of the objects 230 preferably contains information about an individual data  
8 element of the cache 42 of Figure 1. In one embodiment, each object 230 stores a header 230  
9 identifying the data element modeled by the object 230. Each object 230 may also store a  
10 history 234 of the represented data element. In one embodiment, the history 234 is  
11 represented by a priority value 236 assigned to the data element. A marker 238 indicating  
12 whether the data element 223 was stored as a result of a prefetch operation or not may also  
13 be stored within the object 230. A time stamp 240 indicating when a data element first  
14 entered the cache is also preferably present within each object 230. Of course, other data that  
15 may be needed to accurately model each data element that resides within the cache 42 of  
16 Figure 1 may likewise be stored in or with the objects 230.  
17  
18  
19  
20

21 In a most basic embodiment, one remote model 220 of the operation of the cache is  
22 maintained within the remote modeling module 210. The remote model 220 preferably  
23 models the contents and operation of the cache. The model 220 is preferably remote to the  
24 cache 42 in that it is not physically within the cache 42 and is preferably also not logically  
25 part of the original cache 42 which, as stated, may be a preexisting and internally unmodified  
26 LRU cache. The cache may be thought of in one embodiment as a "black box," the contents  
of which are remotely modeled, but which is external to the prefetch module 200.

1 The objects 230 modeled as being within the cache according to the model 220 are  
2 linked to the remote model 220 with a data structure, such as a pointer. Within the model  
3 220 or associated with the model 220 is a calculated single reference residency time (SRRT)  
4 225 for the cache. The SRRT 225 is used to determine when an object has fallen out of the  
5 cache. For instance, in one embodiment, when the time stamp 240 for an object 230 is found  
6 to exceed the SRRT of the cache 42, that object 230 is removed from the model 220 of the  
7 cache. A counter 222 may also be present within the model 220 to store hit rates, as will be  
8 described below.  
9

11 The interface module 202 is shown provided with an intercept module 205. The  
12 intercept module 205 is preferably programmed or otherwise configured to receive either  
13 periodic statistical data or to receive statistical data periodically from the cache 42 of Figure  
14 1. This data may comprise the hit rate of the cache 42 of Figure 1, the cache size of the cache  
15 42 of Figure 1, and the Input/Output (I/O) rate of the cache 42 of Figure 1. This data is then  
16 passed to a SRRT module 215 of the calculation module 204. The SRRT module 215  
17 calculates the SRRT 225 For the model 220. The calculation module 204 is also preferably  
18 provided with a comparison module 212 that is programmed or otherwise configured to  
19 receive the priority value 236 of a data element and to compare that priority value 236 to a  
20 threshold value 214 to determine whether the priority value is greater than the threshold value  
21 214. This determination is then used to determine whether to prefetch the data element, as  
22 will be explained.  
23  
24  
25  
26

The intercept module 205 is preferably configured to monitor the transmission of  
data between the host 12 and the cache 42 and to intercept I/O requests. Other I/O requests

1 to the cache 42 from other stations 12 may not be intercepted. The intercepted I/O requests  
2 are then examined for data elements requested to be transferred. Briefly, when such a data  
3 element is requested, the calculation module confers with the remote modeling module 210  
4 to determine how many logically adjacent preceding data elements are present within the  
5 cache, and if sufficient preceding data elements are present, notifies the prefetch request  
6 module 208 to schedule a prefetch of the next  $n$  logically successive data elements. The  
7 variable  $n$  is the numeral 1 in one embodiment, but could be any number and is in one  
8 embodiment experimentally determined for optimal operation of the cache 42.  
9

11 In one embodiment, the prefetch request module 208 requests the prefetch by issuing  
12 an artificial I/O command that is not solicited by the host processor 16 and as such is ignored  
13 by the processor 16. The cache 42, however, is “fooled” into believing that a real I/O request  
14 has been made and fetches the data element from the storage device 45, keeping a copy in the  
15 cache 42.  
16

18 In one embodiment, the comparison module 212 compares a priority value 236 of  
19 the requested data element and compares it to a threshold value 214 to make the  
20 determination whether it is likely that successive data elements will be requested by the host  
21 12. The priority value 236 is, in one embodiment, calculated according to the number of  
22 logically adjacent preceding data elements that are present in the model 220 of the cache. The  
23 threshold value is first calculated and then continually optimized by the dynamic threshold  
24 optimization module 206 in a manner that will be described below.  
25  
26

In a second, more preferred embodiment, a plurality of models 220 are maintained within the remote modeling module 210. These models in one embodiment comprise a minus

1 model 224, a baseline model 226, and a plus 1 model 228. In the baseline model 226, the operation of the cache 42 is modeled using the current threshold value 214. In the minus 1 model 224, the cache is modeled using the current threshold value minus one whole value. In the plus 1 model, the cache is modeled using the current threshold value plus one. Periodically, as will be explained, the dynamic threshold optimization module 206 compares the operation of the three models 224, 226, 228 and updates the threshold value 214 according to which of the three values is at the time providing the most optimal performance of the cache.

Figure 3 is a schematic flow chart diagram illustrating one embodiment of a data prefetch scheduling method 300 of the present invention. The method 300 in one embodiment may be considered to be the method of use of the prefetch module 200 of Figure 2, but may also be used independent of the prefetch module 200 of Figure 2. In one embodiment, the method 300 operates with three concurrent operations, denoted at 301, 309, and 319. The method 300, may upon initialization, step through each step 302 - 334 in sequence, and after that, the separate operations may loop separately.

The data prefetch scheduling method 300 preferably remotely models the performance of the LRU cache 42 of Figure 1, and in so doing the interface module preferably periodically retrieves performance and statistical data from the LRU cache 42 of Figure 1. In a step 302, the size of the LRU cache 42 is determined and is preferably measured in units of I/O requests. The size of the LRU cache 42 is preferably fetched from the LRU cache 42 through the interface module 202, but may be determined in other manners such as preprogramming or direct user entry. A step 301 indicates that after the first iteration of the



1 steps 302 - 308, the operation 309 loops at intervals such as every x minutes, where x may  
2 be any number and is in one embodiment one minute. Of course, if the cache size is static,  
3 step 302 need be conducted only one time upon startup.  
4

5 The Input/Output (I/O) rate of the cache 42 of Figure 1 is fetched from the cache 42  
6 of Figure 1 by the interface module 204 of Figure 2 in a step 304. The I/O rate refers to the  
7 number of I/O requests sent to the cache 42 of Figure 1 by one or more computer stations 12  
8 of Figure 1 in a given period of time. The I/O rate includes both hits and "misses" in the  
9 cache. The hit rate of the cache 42 of Figure 1 is also preferably fetched from the cache 42  
10 of Figure 1 by the interface module 202 of Figure 2 in a step 306. The hit rate refers to the  
11 the number of I/O requests, i.e., "hits," that were found to reside in the cache at the time of  
12 the I/O request over a given period of time. The hit rate may also be used to determine a miss  
13 ratio for the cache where the miss ratio is one minus the hit rate divided by the number of I/O  
14 requests.  
15  
16  
17

18 The hit rate, in the preferred embodiment, is calculated based upon a logarithmic rate  
19 of increase with increased cache size. Calculating the hit rate from the cache size in this  
20 manner is necessary, because hit rates do not typically increase in a linear manner with an  
21 increasing cache size. Thus, a logarithmic rate of increase is considered a close  
22 approximation of the hit rate v. cache size, where the hit rate is estimated to increase at an  
23 increasingly lower rate with the increase in cache size.  
24  
25

26 Once the cache size, I/O rate, and hit rate are fetched, the calculation module 204 of  
Figure 2 preferably uses this data to calculate the SRRT 216, which is an approximation of  
the average time it takes a data element that has been fetched into the cache to "fall through"

1 or be removed from the cache 42 of Figure 1. The SRRT 216 is later used at a step 322 to  
2 determine which data elements have “fallen out of the cache.”  
3

4 In a preferred embodiment, the SRRT is calculated by estimating the hit rate for all  
5 cache sizes 1 to  $n$ , where  $n$  is the size of the cache as determined in the step 302, and is  
6 adjusted downward by the portion of the cache assumed to be holding prefetched requests.  
7 Then, using the estimated hit rate of the different cache sizes, the miss ratio for all cache sizes  
8 1 to  $n$  is calculated. This is done by assuming that the effectiveness of the cache decays  
9 exponentially in relation to the size of the cache. Subsequently, using the I/O rate of the  
10 cache, the miss ratio of the cache, and the expected miss ratio for the cache one size smaller,  
11 the SRRT is calculated using Formula 1,  
12  
13

$$\frac{1}{r} \left( 1 + p + \frac{n-1}{m(n-1)} \right)$$

14  
15  
16  
17  
18 Formula 1

19 where  $r$  is the I/O rate of the cache,  $p$  is the number of prefetches conducted by the  
20 prefetching algorithm,  $n$  is the size of the cache as determined in the step 302, and  $m[i]$  is the  
21 expected miss ratio of the cache when the cache is of size  $i$ . Of course, other methods exist  
22 for calculating the SRRT for an LRU cache and may be used in place of Formula 1. The steps  
23 304, 306 and 308 of the data prefetch scheduling method 300 are preferably performed  
24 periodically as indicated by the step 301. In the embodiment described above, where a  
25 plurality of model 220 is maintained, the SRRT that would occur under operation of each  
26

1 model 220 is calculated. Thus, in the depicted embodiment of Figure 2, three SRRT values  
2 are calculated at the step 308 for every iteration of the operation 301.

3  
4 The second operation 309, consisting of the steps 310 - 318, is initially conducted  
5 after the operation 301 and then at intervals as determined by the step 310. In one  
6 embodiment, the operation 309 is conducted upon the occurrence of every  $y$  intercepted I/O  
7 requests, where  $y$  may be determined experimentally. In one embodiment,  $y$  comprises  
8 10,000. In the initial iteration of the operation 309, as data may not exist, a preselected  
9 threshold value may be used and the steps 310 - 318 skipped. Otherwise, stored data may be  
10 used.  
11

12  
13 At a step 312, the expected (i.e., modeled) hit rates of prefetched data of the various  
14 models 220 (i.e., 224, 226, 228) are compared. The results may then optionally undergo  
15 weighting or other conditioning at a step 314 to determine trends so that the threshold value  
16 214 does not change too rapidly. At a step 316, the various expected hit rates are assessed  
17 to determine whether there is a more optimal threshold valve. If one of the models other than  
18 the baseline model 226 has a higher expected hit rate, the value (e.g., threshold minus 1 or  
19 threshold value plus 1) corresponding to that model then affects the weighting process, which  
20 determines the threshold value to be used as the new current threshold value 214. In one  
21 embodiment, the hit rate that are to be compared are stored as counters 222 within each of  
22 the models 220 and are updated at a step 324 as will be described.  
23  
24  
25

26 Figure 5 illustrates in greater detail, one embodiment of a method for conducting the  
threshold value update operation 309 of Figure 3. The method of figure 5 begins at step 400,  
and is conducted every  $y$  I/O request, as indicted by step 310 on figure 3. At step 401, the

1 hit rate counter 222 of the baseline model 226 of Figure 2 is retrieved. At a step 402, the  
2 baseline model counter is compared with the hit rate counter 222 of the minus 1 model 224.  
3  
4 At a step 403, the results are weighted. That is, the results may be averaged over every  $m$   
5 iterations of the operation 309. At a step 404, the results are observed. Thus, after the  
6 weighting, if the baseline counter is determined to be less than the minus 1 counter, the  
7 dynamic threshold is decremented by one integer value at a step 406. The operation 309 then  
8 returns to step 310 as indicated at 407. If the results of the observation of step 404 are  
9 negative, the baseline counter is compared with the hit rate counter of the plus one model 228  
10 at a step 408. Similarly, the results are weighted at a step 409. If, in the weighted results,  
11 the baseline counter is less than the plus 1 counter, the dynamic threshold is incremented as  
12 indicated at a step 412, and the operation 309 returns back to the step 310 as indicated at  
13 413. If the result of the query of step 10 is no, the operation 309 also returns directly to step  
14 310 as indicated at 414.  
15  
16  
17

18 Once a threshold value has been selected and/or optimized, the method 300 proceeds  
19 to the third operation 319. The operation 319 is also initially conducted in sequence with the  
20 operations 301 and 309 and then is conducted independently as data requests are intercepted.  
21  
22 At a step 320, the operation 319 waits for the intercept module 205 to intercept an I/O  
23 request. When an I/O request is intercepted, a data element requested from the storage  
24 device 45 or other source is extracted. The models 220 are then updated at a step 322. This  
25 preferably comprises comparing the timestamps of the objects 230 within each of the models  
26 220 to determine which have exceeded the SRRT for that model. Those exceeding the SRRT  
are removed from the model as being likely to have been removed from the cache 42 under

1 that model. Of course, the step 322 could be conducted at any time, and could be a separate  
2 independent operation. In one embodiment, this step is repeated later as indicated by step  
3 334.  
4

5 At a step 324, the hit rates are updated. In one embodiment, this comprises  
6 examining the data element identified at the step 320 and comparing that data element to each  
7 of the models 220. If the data element is within a model and is marked with the marker 238  
8 as having been prefetched, the counter 222 of that model is incremented.  
9

10 At a step 326, the models 224, 226, 228 are examined to see if the data element  
11 logically preceding the intercepted data element is present within each model. At a step 328,  
12 it is determined whether a prefetch of the successive data element(s) is to be conducted  
13 according to each of the models 220 and the models 220 are separately updated according  
14 to the results of the determination. One embodiment of a method of achieving this process  
15 is illustrated in greater detail in Figure 4.  
16  
17

18 Referring to Figure 4, at a step 350 the method begins, and as indicated at a step 352,  
19 is conducted for each of the models 220, beginning with the minus 1 model 224. At a step  
20 354, the data element is received into the prefetch module as an object 230. At a step 355,  
21 an object 230 is created or updated, depending on whether the data element is already in the  
22 cache. In so doing, a timestamp 240 is applied and a header 232 is assigned. The marker 238  
23 is also set to indicate whether the data element was prefetched or not. The data element is  
24 modeled as being the youngest element in the cache 42, as indicated at a step 356, in  
25 accordance with how the cache 42, being a LRU cache, literally treats the data element.  
26

At a step 358, the prefetch module 260 determines whether the logically preceding data element is present in the cache 42. If so, at a step 360, the data element is assigned a value for the respective model, which in one embodiment comprises the priority value 236 of the logically preceding element plus one (optionally up to some maximum value selected to prevent overflow). Because each preceding element was similarly updated, the priority value 236 thus represents the amount of preceding data elements stored in the model. Thus, the object 230 corresponding to the data element may be given a separate value 236a, 236b, 236c, for each of the models 220 as the method 238 is conducted for each of the models 220. If the immediately preceding data element is not present in the cache, at a step 360, a standard initial priority value is assigned to the object 230 corresponding to the intercepted data element. In one embodiment, this value is 1.

Thereafter, at a step 364, the assigned priority value of the data element is compared to the threshold value of the model for which the iteration is being conducted, preferably by the calculation module 204. At a step 366, the calculation module 204 determines whether a prefetch is to be scheduled. If the priority value 236 of the object is equal to or greater than the threshold value 214, a prefetch is considered to have been scheduled and the model is updated accordingly at a step 366. Also at the step 366, the size of the cache in each model 220 is updated. This preferably comprises subtracting the number of prefetched elements from the cache size of step 302 and using the new cache size for future calculations of the SRRT for the model.

Returning to Figure 3, at a step 330, the operation 319 determines whether the prefetch of the data element is to be conducted according to the baseline model. That is, is

1 the priority valve assigned to the data element greater than the threshold valve? If so, at a  
2 step 332, the prefetch is scheduled. Scheduling the prefetch preferably comprises passing an  
3 unsolicited I/O request to the cache 42, as discussed above.  
4

5 In accordance with the present invention as described above, it should be apparent  
6 that sequential prefetching of data in a preexisting and/or dedicated LRU cache can now be  
7 conducted without significant likelihood of the prefetched data falling out of the cache before  
8 being re-referenced. Additionally, this can be accomplished with minimum interference to  
9 the "normal" operation of the cache and without having to internally modify the LRU cache  
10 or the native operation of the LRU cache. Accordingly, such caches can be enhanced for  
11 greater performance with the use of the method and system of the present invention merely  
12 with, in one embodiment, the addition of software to a host or between the cache and the  
13 host.  
14  
15

16 The present invention may be embodied in other specific forms without departing  
17 from its spirit or essential characteristics. The described embodiments are to be considered  
18 in all respects only as illustrative and not restrictive. The scope of the invention is, therefore,  
19 indicated by the appended claims rather than by the foregoing description. All changes which  
20 come within the meaning and range of equivalency of the claims are to be embraced within  
21 their scope.  
22  
23

24 What is claimed is:  
25  
26





1           5. The method of claim 2, wherein determining whether to schedule a prefetch of  
2 data into the cache further comprises checking the model to determine whether the requested  
3 data element is likely to be present within the cache.  
4

5  
6           6. The method of claim 1, wherein determining whether to schedule a prefetch  
7 further comprises examining the history of a second data element stored logically adjacent to  
8 the requested data element in the storage device.  
9

10  
11           7. The method of claim 1, wherein remotely modeling the cache further comprises:  
12  
13                 determining the size of the cache;  
14                 periodically fetching the I/O rate of the cache; and  
15                 periodically fetching the hit rate of the cache.  
16

17  
18           8. The method of claim 1, wherein remotely modeling the cache further comprises  
19 periodically calculating a single reference residency time (SRRT) for a data element within  
20 the cache.  
21

22  
23           9. The method of claim 1, wherein remotely modeling the cache further comprises  
24 the step of treating a requested data element as the youngest member of the cache when the  
25 requested data element is already present in the cache.  
26

1           10. The method of claim 1, wherein remotely modeling the cache further comprises  
2 determining whether the data element preceding the requested data element in a sequential  
3 stream of data is also present in the cache.  
4

5  
6           11. The method of claim 1, wherein remotely modeling the cache further comprises  
7 assigning a priority value to the requested data element.  
8

9  
10           12. The method of claim 11, wherein assigning a priority value further comprises  
11 assigning a priority value comprising the priority value assigned to the preceding data element  
12 plus one when the preceding data element is found to be present in the cache.  
13

14  
15           13. The method of claim 11, wherein determining whether to schedule a prefetch of  
16 a data element further comprises comparing the priority value of the requested element with  
17 a dynamic threshold.  
18

19  
20           14. The method of claim 13, further comprising prefetching the requested data  
21 element into the cache if the priority value of the requested data element is greater than the  
22 dynamic threshold.  
23

24  
25           15. The prefetch method of claim 1, further comprising periodically reevaluating the  
26 performance of the cache model.

1           16. The method of claim 15, wherein periodically reevaluating the performance of  
2 the cache further comprises determining if the dynamic threshold used in the internal model  
3 of the cache accurately models the performance of the cache.  
4

5  
6           17. The method of claim 16, wherein determining if the dynamic threshold  
7 accurately models the performance of the cache comprises comparing the performance of the  
8 dynamic threshold with an alternate dynamic threshold.  
9

10  
11           18. The method of claim 15, further comprising automatically updating the dynamic  
12 threshold used in the internal model of the cache when another dynamic threshold is deemed  
13 to be more effective.  
14

15  
16           19. The method of claim 1, wherein making a cache management decision  
17 comprises deciding to schedule a prefetch, and further comprising scheduling a prefetch by  
18 sending an I/O request to the cache.  
19  
20

21  
22           20. A method for scheduling prefetches in a data storage system having a host and  
23 a cache, the method comprising the steps of:  
24

25           providing a cache for caching Input/Output (I/O) data;

26           providing a prefetch module remote to the cache;

1 remotely modeling the cache within the prefetch module and determining whether  
2 to schedule a prefetch of data into the cache according to the results of the step of remotely  
3 modeling the cache, the step of remotely modeling the cache module further comprising:

4 examining the history of a data element in the cache;

5 assigning a priority value to the data element according to its history;

6 comparing that priority value to a predetermined threshold value;

7 determining the size of memory used in the cache;

8 periodically fetching the I/O rate of the cache from the cache;

9 periodically fetching the hit rate of the cache from the cache; and

10 determining a single reference residency time for a data element within the

11 cache;

12 intercepting a stream of I/O information from the host to the cache to locate a  
13 requested data element;

14 determining if the requested data element in the stream of I/O information is already  
15 present within the cache;

16 making the requested data element the youngest member of the cache;

17 determining if the data element preceding the requested data element is present in the  
18 cache;

19 assigning a priority value to the requested data element;

20 periodically reevaluating the performance of the cache versus an internal model of  
21 the cache if the number of I/O requests received by the cache is greater than a predetermined  
22 number;

1 updating the dynamic threshold used in the internal model of the cache if the dynamic  
2 threshold value does not adequately model the performance of the cache;

3 comparing the priority value of the requested data element with the dynamic  
4 threshold value; and

5 prefetching the requested data element if the priority value of the requested data  
6 element is greater than the dynamic threshold value by passing an I/O request of the data  
7 element to the cache.  
8  
9

10  
11 21. A data prefetch scheduling system comprising:

12 a cache configured to communicate with a host; and

13 a remote prefetch module configured to communicate with the host and the cache  
14 and configured to determine whether to schedule a prefetch of data into the cache; and

15 a modeling module operating within the prefetch scheduling module configured to  
16 model the cache.  
17  
18

19  
20 22. The data prefetch scheduling system of claim 21, wherein the cache comprises  
21 a least recently used (LRU) cache.  
22

23  
24 23. The data prefetch scheduling system of claim 22, wherein the LRU cache is a  
25 native LRU-only cache that is not internally modified.  
26

1           24. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
2 module further comprises a calculation module configured to compare a priority value  
3 assigned to a data element to a threshold value and determine whether to schedule a prefetch  
4 of the data element.  
5

6  
7           25. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
8 module further comprises a dynamic threshold optimization configured to calculate and  
9 update a dynamic threshold used in determining whether to prefetch data.  
10

11  
12           26. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
13 module is configured to model the cache for use in determining when to prefetch I/O data into  
14 the cache.  
15

16  
17           27. The data prefetch scheduling system of claim 21, wherein the remote prefetch  
18 module is configured to prefetch data into the cache according to a priority scheme that takes  
19 into account the run length of each sequential I/O stream.  
20

21  
22           28. The data prefetch scheduling system of claim 21, further comprising a prefetch  
23 request module, the prefetch request module configured to request a data I/O from the cache  
24 when the remote prefetch module determines that a prefetch is to be conducted.  
25  
26

Case 1:12-cv-00000

BRIAN C. KUNZLER  
ATTORNEY AT LAW  
10 WEST 100 SOUTH, SUITE 425  
SALT LAKE CITY, UTAH 84101

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26

SCHEDULING SEQUENTIAL DATA PREFETCHES IN A  
PREEXISTING LRU CACHE

ABSTRACT OF THE INVENTION

A shared system memory, such as a cache, buffers Input/Output (I/O) requests between one or more host computers and one or more data storage servers or devices. The cache may be configured to operate natively as a least-recently-used (LRU)-only cache and may be optimized for random data accesses. Data buffered by the cache may be part of a sequential data stream for which prefetching data is desirable. A remote prefetch module is provided between the cache and the host to conduct prefetching without internally modifying the cache. The remote prefetch module maintains a model of the cache. Using the model, the prefetch module anticipates whether data is likely to be part of a sequential steam of data passed between a host and a data storage device. If so, the prefetch module schedules a prefetch of the data. The prefetch may be achieved by sending an I/O request to the data server or device. The remote cache model minimizes impacts to random access data hits by minimizing the likelihood of prefetching data which is not used and further enhances the efficiency of the successful identification of likely prefetch candidates.

E:\Work\Client Files\1200 SanJose\1211\1200 2 11 pap




Fig. 1






Fig. 2

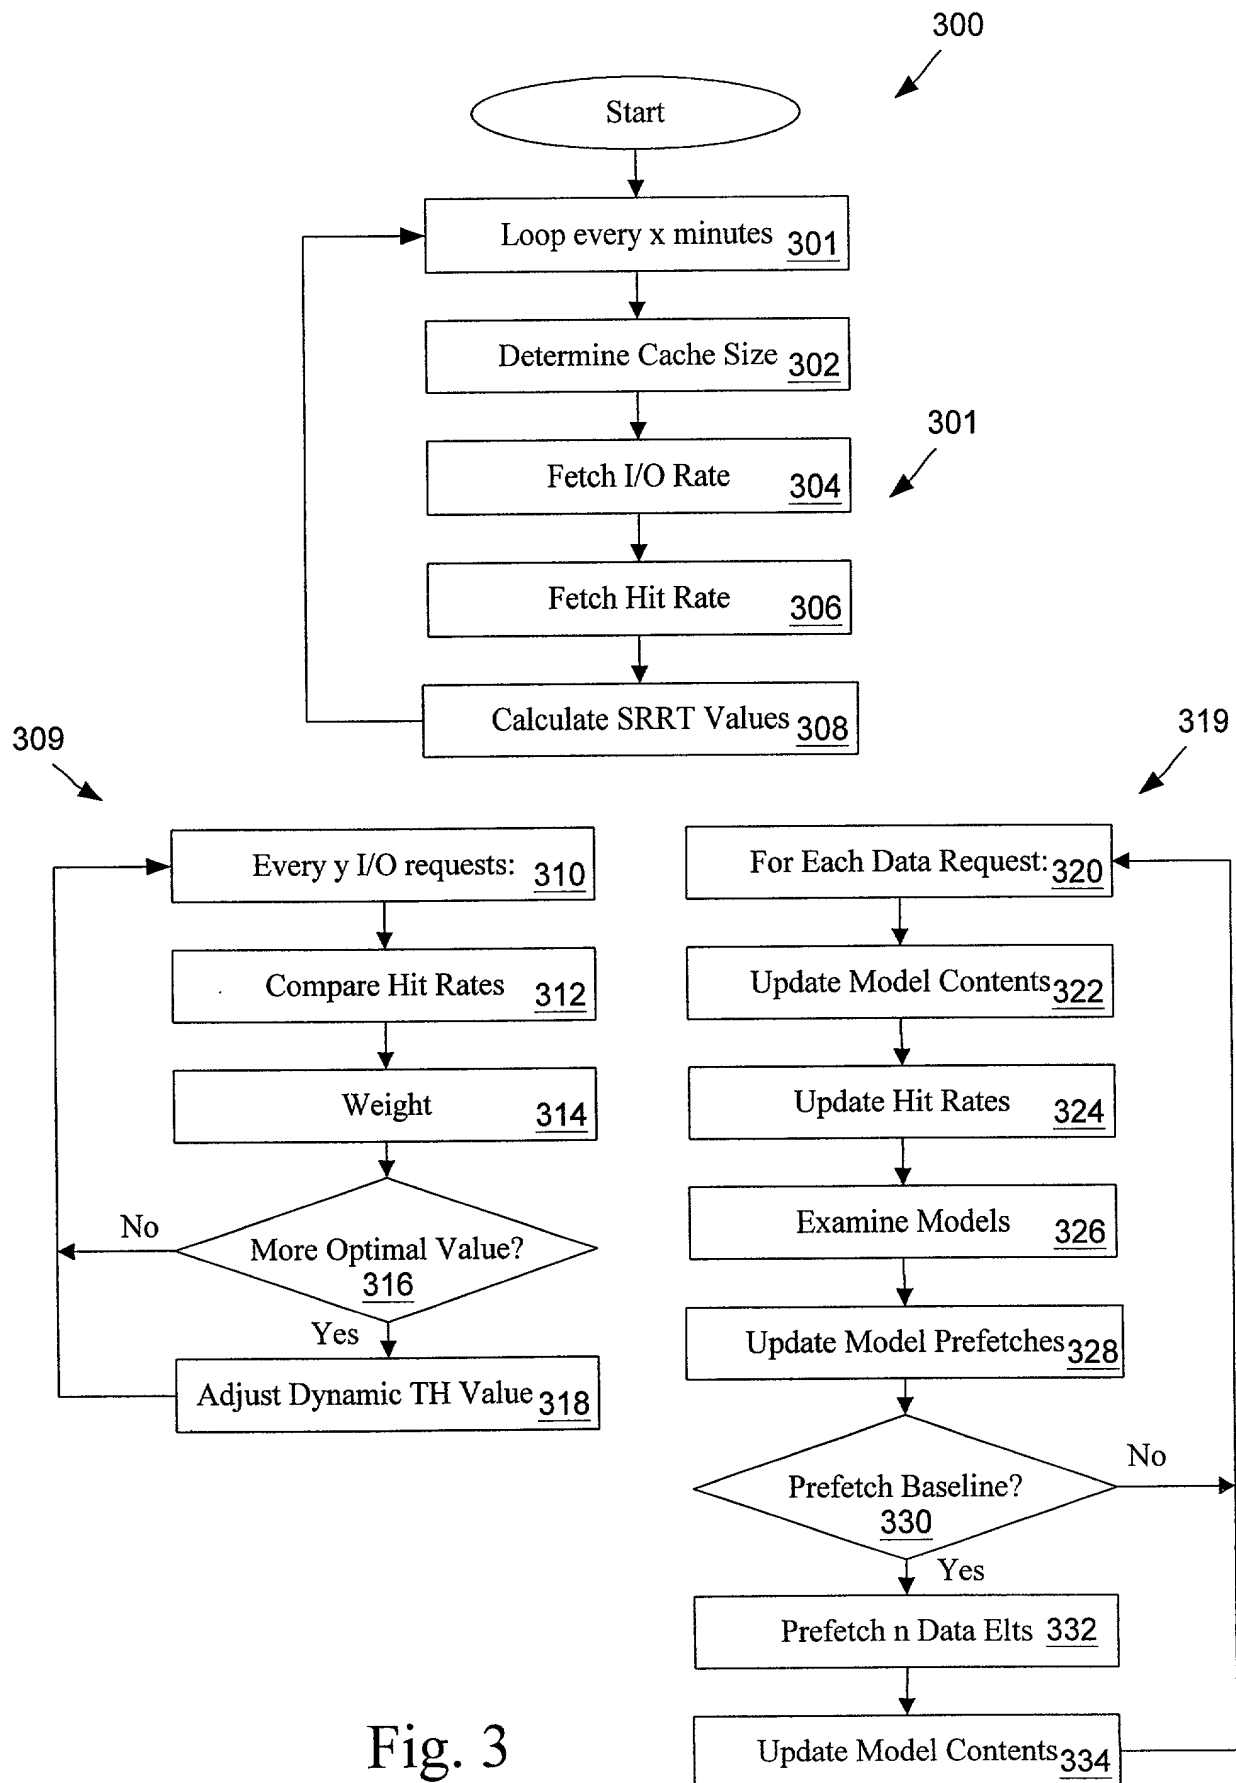


Fig. 3

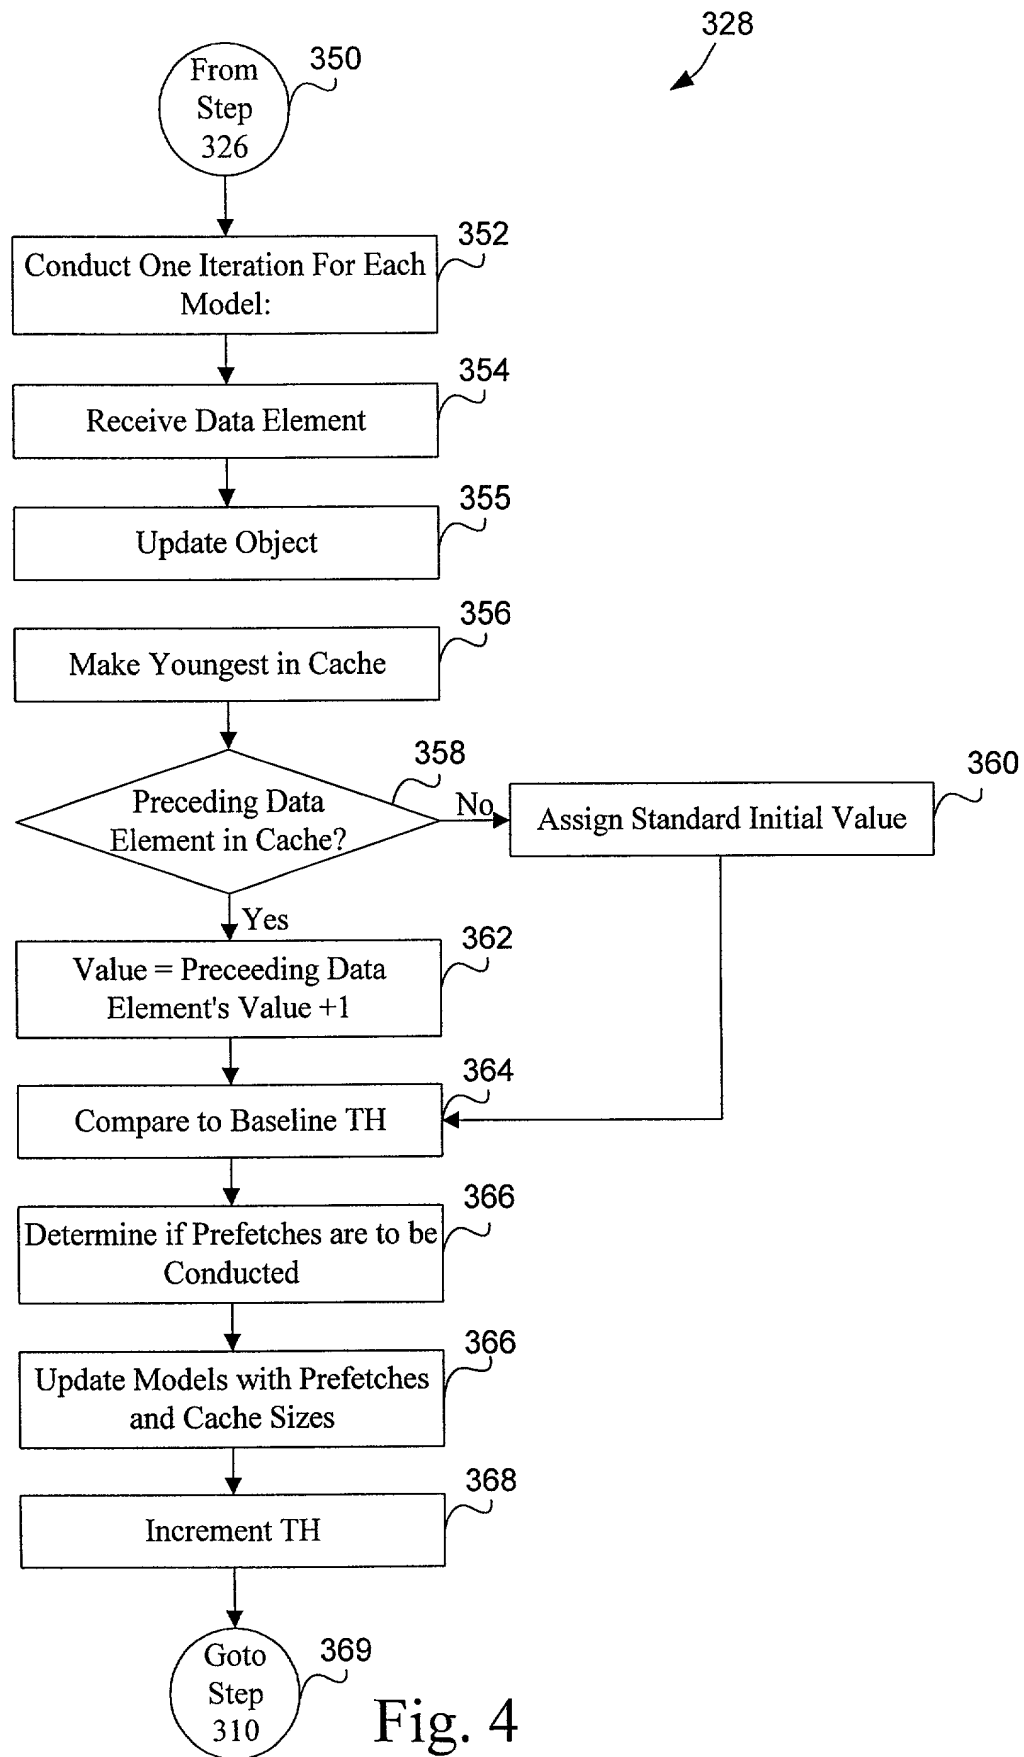


Fig. 4

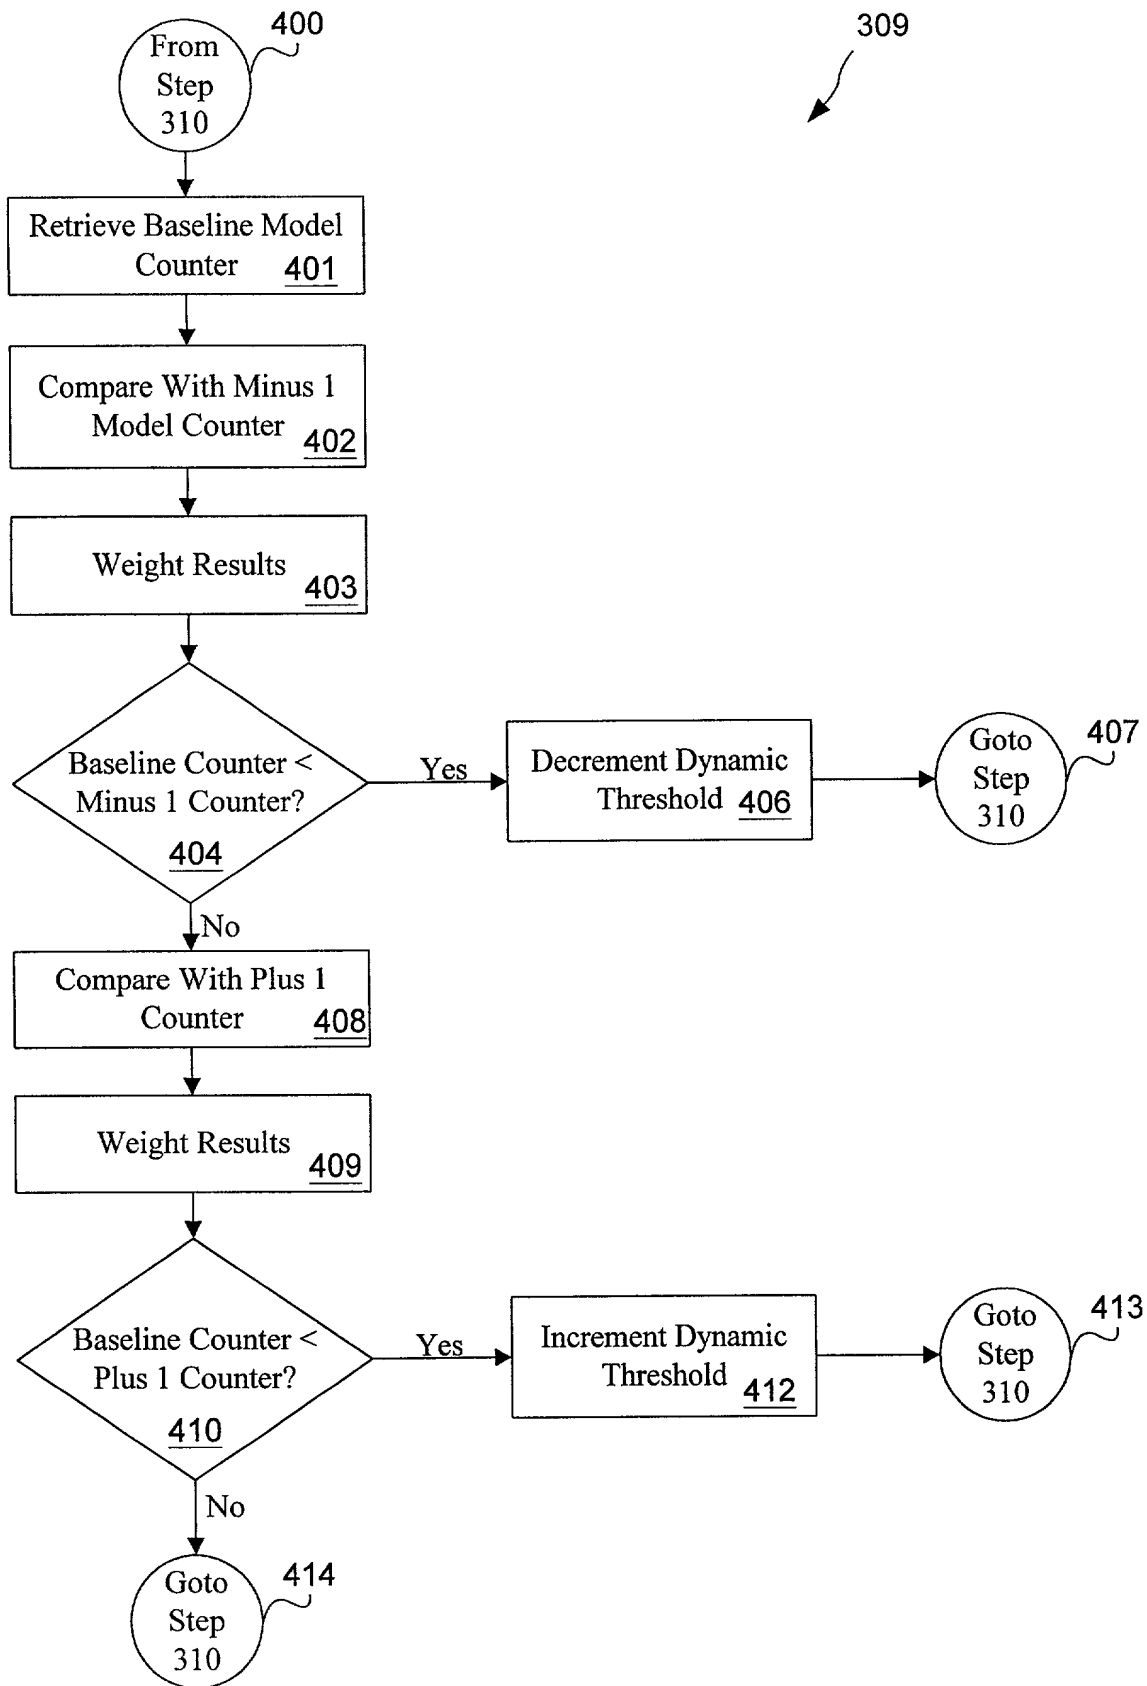


Fig. 5

**DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence and citizenship are as stated below next to my name;

I believe I am the original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled

**PRESCHEDULING SEQUENTIAL DATA PREFETCHES IN A PREEXISTING LRU CACHE**

the specification of which (check one)

  X   is attached hereto.  
       was filed on \_\_\_\_\_  
as Application Serial No. \_\_\_\_\_  
and was amended on \_\_\_\_\_ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)			Priority Claimed
<u>  none  </u>	<u>                    </u>	<u>                    </u>	<u>      </u> Yes <u>      </u> No
(Number)	(Country)	(Day/Month/Year filed)	

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56, which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

<u>  none  </u>	<u>                    </u>	<u>                    </u>
(Application Serial No.)	(Filing Date)	(Status) (patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

**POWER OF ATTORNEY:** As named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (list name and registration number)

Esther E. Klein	34,337	Douglas R. Millett	31,784
Paik Saber	37,494	Abdy Raissinia	38,686
Edward A. Pennington	32,588	Christopher A. Hughes	26,914
Joseph C. Redmond, Jr.	18,753	John E. Hoel	26,279
Robert B. Martin	26,945	Brian C. Kunzler	38,527
Randall J. Bluestone	40,518		

Send correspondence to:

Brian C. Kunzler  
10 West 100 South Suite 425  
Salt Lake City, Utah 84101  
Telephone: (801) 994-4646

Full name of sole or first joint-inventor: Kevin Frank Smith

Inventor's signature: Kevin Frank Smith Date: 10/10/2000

Residence: 17941 Oak Grove Drive, Morgan Hill, California 95037-4224

Citizenship: United States of America

Post Office Address: Same